

Why Bother With Syntax?

Joseph Y. Halpern*
 Cornell University
 Computer Science Department
 Ithaca, NY 14853
 halpern@cs.cornell.edu
<http://www.cs.cornell.edu/home/halpern>

June 18, 2015

Rohit and I go back a long way. We started talking about Dynamic Logic back when I was a graduate student, when we would meet at seminars at MIT (my advisor Albert Meyer was at MIT, although I was at Harvard, and Rohit was then at Boston University). Right from the beginning I appreciated Rohit’s breadth, his quick insights, his wit, and his welcoming and gracious style. Rohit has been interested in the interplay between logic, philosophy, and language ever since I’ve known him. Over the years, both of us have gotten interested in game theory. I would like to dedicate this short note, which discusses issues at the intersection of all these areas, to him.¹

1 Introduction

When economists represent and reason about knowledge, they typically do so at a *semantic* (or set-theoretic) level. Events correspond to sets and the knowledge operator maps sets to sets. On the other hand, in the literature on reasoning about knowledge in philosophy or logic, there is an extra layer

*Supported in part by NSF grants IIS-0812045, IIS-0911036, and CCF-1214844, AFOSR grants FA9550-08-1-0266 and FA9550-12-1-0040, and ARO grant W911NF-09-1-0281.

¹Some material in this note appears in [Halpern 2003, Section 7.2.4].

of what may be viewed as unnecessary overhead: syntax. There have been papers in the economics literature that have argued for the importance of syntax. For example, Feinberg [2000, p. 128] says “The syntactic formalism is the more fundamental and—intuitively—the more descriptive way to model economic situations that involve knowledge and belief . . . It is fine to use the semantic formalism, as long as what we say semantically has a fairly clear intuitive meaning—that it can be said *in words*. This amounts to saying that it can be stated syntactically.”

In this brief note, I point out some technical advantages of using syntax. Roughly speaking, they are the following:

1. Syntax allows us to make finer distinctions than semantics; a set may be represented by the two different expressions, and an agent may not react to these expressions in the same way. Moreover, different agents may react differently to the same expression, that is, the expression may represent different sets according to different agents.
2. Syntax allows us to describe in a model-independent way notions such as “rationality”. This enables us to identify corresponding events (such as “agent 1 is rational”) in two different systems.
3. The structure of the syntax provides ways to reason and carry out proofs. For example, many technical results proceed by induction on the structure of formulas. Similarly, formal axiomatic reasoning typically takes advantage of the syntactic structure of formulas.

In the rest of this note, I briefly review the semantic and syntactic approaches, explain the advantages listed above in more detail, and point out ways in which the economics literature has not exploited the full power of the syntactic approach thus far. In another paper [Halpern 1999], I have made essentially the opposite argument, namely, that computer scientists and logicians have not exploited the full power of the semantic approach, pointing out that we can often dispense with the overhead of syntax and the need to define a \models relation by working directly with the semantics. I believe both arguments! This just goes to show that each community can learn from the approaches used by the other community. I return to this point in the concluding section of this paper.

2 Semantics vs. Syntax: A Review

The standard approach in the economic literature on knowledge starts with what is called in [Fagin, Halpern, Moses, and Vardi 1995] an *Aumann structure* $A = (\Omega, \mathcal{P}_1, \dots, \mathcal{P}_n)$, where Ω is a set of states of the world and \mathcal{P}_i , $i = 1, \dots, n$ are partitions, one corresponding to each agent. Intuitively, worlds in the same cell of partition i are indistinguishable to agent i . Given a world $\omega \in \Omega$, let $\mathcal{P}_i(\omega)$ be the cell of \mathcal{P}_i that contains ω .

Knowledge operators K_i , $i = 1, \dots, n$ mapping events in (subsets of) Ω to events are defined as follows:

$$K_i(A) = \{\omega : \mathcal{P}_i(\omega) \subseteq A\}.$$

We read $K_i(A)$ as “agent i knows A ”. Intuitively, it includes precisely those worlds ω such that A holds at all the worlds that agent i cannot distinguish from ω .

The logical/philosophical approach adds an extra level of indirection to the set-theoretic approach. The first step is to define a *language* for reasoning about knowledge, that is, a set of well-formed formulas. We then associate an event with each formula. We proceed as follows.

We start with a set Φ of primitive propositions p_1, p_2, \dots representing propositions of interest. For example, p_1 might represent “agent 1 is rational” and p_2 might represent “agent 2 is following the strategy of always defecting (in a game of repeated Prisoner’s Dilemma)”. We then close off this set under conjunction, negation, and the *modal operators* K_1, \dots, K_n . Thus, if φ and ψ are formulas, then so are $\varphi \wedge \psi$, $\neg\varphi$, and $K_i\varphi$, for $i = 1, \dots, n$. We typically write $\varphi \vee \psi$ as an abbreviation for $\neg(\neg\varphi \wedge \neg\psi)$ and $\varphi \Rightarrow \psi$ as an abbreviation for $\neg\varphi \vee \psi$. Thus, we can write formulas such as $K_1 K_2 p_1 \wedge \neg K_2 K_1 p_2$, which could be read as “agent 1 knows that agent 2 knows that agent 1 is rational and agent 2 does not know that agent 1 knows that agent 2 is following the strategy of always defecting”. It should be stressed that a formula like $K_1 K_2 p_1 \wedge \neg K_2 K_1 p_2$ is just a string of symbols, not a set.

A *Kripke structure* is a tuple $M = (\Omega, \mathcal{P}_1, \dots, \mathcal{P}_n, \pi)$, where $(\Omega, \mathcal{P}_1, \dots, \mathcal{P}_n)$ is an Aumann structure, and π is an *interpretation*, that associates with each primitive proposition an event in Ω .² We can then extend π inductively to associate with each formula an event as follows:

²In the literature, π is often taken to associate a *truth value* with each primitive proposition at each world; that is, $\pi : \Phi \times \Omega \rightarrow \{\text{true}, \text{false}\}$. Using this approach, we can then associate with each primitive proposition p the event $\{\omega : \pi(p, \omega) = \text{true}\}$. Conversely,

- $\pi(\varphi \wedge \psi) = \pi(\varphi) \cap \pi(\psi)$
- $\pi(\neg\varphi) = \overline{\pi(\varphi)}$ (where \overline{E} denotes the complement of E)
- $\pi(K_i\varphi) = \mathsf{K}_i(\pi(\varphi))$.³

Notice that not every subset of Ω is necessarily definable by a formula (even if Ω is finite). That is, for a given subset $E \subseteq \Omega$, there may be no formula φ such that $\pi(\varphi) = E$. The set of events definable by formulas form an algebra.

If all that is done with a formula is to translate it to an event, why bother with the overhead of formulas? Would it not just be simpler to dispense with formulas and interpretations, and work directly with events? It is true that often there is no particular advantage in working with syntax. However, sometimes it does come in handy. Here I discuss the three advantages mentioned above in some more detail:

1. The two events E and $(E \cap F) \cup (E \cap \overline{F})$ denote the same set. Hence, so do $\mathsf{K}_i(E)$ and $\mathsf{K}_i((E \cap F) \cup (E \cap \overline{F}))$. Using the set-theoretic approach, there is no way to distinguish these events. Even if we modify the definition of K , and move to non-partitional definitions of knowledge [Bacharach 1985; Fagin, Halpern, Moses, and Vardi 1995; Geanakoplos 1989; Samet 1990], we still cannot distinguish these formulas. In propositional logic, the formula p is equivalent to $(p \wedge q) \vee (p \wedge \neg q)$. (Of course, these formulas were obtained from the events by substituting p for E , q for F , and replacing \cap , \cup , and $-$ by \wedge , \vee , and \neg , respectively.) In the standard semantics defined above, the formulas $K_i p$ and $K_i((p \wedge q) \vee (p \wedge \neg q))$ are also equivalent. However, an agent may not recognize that the formulas p and $(p \wedge q) \vee (p \wedge \neg q)$ are equivalent, and may react differently to them. There are approaches to giving semantics to knowledge formulas that allow us to distinguish these formulas [Fagin, Halpern, Moses, and Vardi 1995, Chapter 9]. This issue becomes important if we are trying to model resource-bounded notions of knowledge. We may, for example, want to restrict the agent's knowledge to formulas that are particularly simple, according to some

given a mapping from primitive propositions to events, we can construct a mapping from $\Phi \times \Omega$ to truth values. Thus, the two approaches are equivalent.

³In the literature, one often sees the notation $(M, \omega) \models \varphi$, which is read as “the formula φ is true at world ω in Kripke structure M ”. The definition of \models recapitulates that just given for π , so that $\omega \in \pi(\varphi)$ iff $(M, \omega) \models \varphi$.

notion of simplicity; thus, it may be the case that $K_i p$ holds, while $K_i((p \wedge q) \vee (p \wedge \neg q))$ does not. Syntactic approaches to dealing with *awareness* [Fagin and Halpern 1988; Halpern and Rêgo 2013] can capture this intuition by allowing agent i to be aware of p but not of $(p \wedge q) \vee (p \wedge \neg q)$. This issue also figures prominently in a recent approach to decision theory [Blume, Easley, and Halpern 2006] that takes an agent’s object of choice to be a syntactic program, which involves tests (which are formulas). Again, an agent’s decision can, in principle, depend on the form of the test. A yet more general approach, where an agent’s utility function is explicitly defined on formulas in the agent’s language, is considered in [Bjorndahl, Halpern, and Pass 2013].

A slightly more general notion of Kripke structure allows us to deal with (at least one form of) *ambiguity*. The standard definition of Kripke structure has a single interpretation π . But, in practice, agents often interpret the same statement differently. What one agent calls “red” might not be “red” to another agent. It is easy to deal with this; we simply have a different interpretation π_i for each agent i (see [Halpern and Kets 2012] for the implications of allowing such ambiguity for game theory).

2. Language allows us to describe notions in a model-independent way. For example, the typical approach to defining rationality is to define the event that agent i is rational in a particular Aumann structure. But suppose we are interested in reasoning about two related Aumann structures, A_1 and A_2 , at the same time. Perhaps each of them has the same set Ω of possible worlds, but different partitions. We then want to discuss “corresponding” events in each structure, for example, events such as “agent 1 is rational” or “agent 2 is following the strategy of always defecting”. This is not so easy to do if we simply use the events. The event “agent 1 is rational” corresponds to different sets in A_1 and A_2 . In the set-based approach, there is no way of relating them. Using syntax, an event such as “agent 1 is rational” would be described by a formula (the exact formula depends, of course, on the definition of rational, a matter of some controversy; see, for example, [Binmore 2009; Blume and Easley 2008; Gilboa 2010]). The same formula may well correspond to different events in the two structures. For example, if the formula involves K_1 , then if different partitions characterize agent 1’s knowledge in A_1 and A_2 , the formula would correspond to different

subsets of Ω . Although the formula describing a notion like “agent 1 is rational” would correspond to different events in different structures, we may well be able to prove general properties of the formula that are true of all events corresponding to the formula. For example, if φ is the formula that says that agent 1 is rational, we may be able to show that $\varphi \Rightarrow K_i \varphi$ is *valid* (true in every state of every Aumann structure); this says that if agent 1 is rational, then agent 1 knows that he is rational.

3. There are standard examples of when the use of formulas can be useful in proofs. For one thing, it allows us to prove results by induction on the structure of formulas. Examples of such proofs can be found in [Fagin, Geanakoplos, Halpern, and Vardi 1992; Fagin, Halpern, Moses, and Vardi 1995]. Secondly, the syntactic structure of a formula we would like to prove can suggest a proof technique. For example, a disjunction can often be proved by cases. Thirdly, when proving that certain axioms completely characterize knowledge, the standard proof involves building a *canonical model*, whose worlds can be identified with sets of formulas [Aumann 1999; Fagin, Halpern, Moses, and Vardi 1995]. This approach is also implicitly taken by Heifetz and Samet [1998]. They construct a *universal type space*, a space that contains all possible types of agents (where a *type*, roughly speaking, is a complete description of an agent’s beliefs about the world and about other agents’ types). Their construction involves certain events that they call *expressions*; these expressions are perhaps best thought of as syntactic expressions in a language of belief.

At some level, the economics community is already aware of the advantages of syntax, and syntactic expressions are used with varying levels of formality in a number of papers. To cite one example, in [Balkenborg and Winter 1997] the notion of an *epistemic expression* is defined, that is, a function from Aumann spaces to events in that Aumann space. An epistemic expression such as R_i , which can be thought of as representing the proposition “agent i is rational” then becomes a mapping from an Aumann structure to the event consisting of all worlds where i is rational.⁴ More complicated epistemic expressions are then allowed, such as $K_1 K_2 (R_3)$. Perhaps not surprisingly, some proofs then proceed by induction on the structure

⁴Not all papers are so careful. For example, Aumann [1995] calls R_i an event. Of course, his intention is quite clear.

of event expressions. It should be clear that event expressions are in fact syntactic expressions, with some overloading of notation. Epistemic expressions are just strings of symbols. The K_i in a syntactic expression is not acting as a function from events to events. It is precisely because they are strings of symbols that we can do induction on their length. Indeed, in [Fagin, Geanakoplos, Halpern, and Vardi 1992], epistemic expressions are introduced in the middle of a proof (under the perhaps more appropriate name *(event) descriptions*, since in fact they are used even to describe non-epistemic events), precisely to allow a proof by induction.

3 Discussion

Expressive power: In the semantic approach, we can apply the knowledge operator to an arbitrary subset of possible worlds. In the syntactic approach, we apply the knowledge operator to formulas. While formulas are associated with events, it is not necessarily the case that every event is definable by a formula. Indeed, one of the major issues of concern to logicians when considering a particular syntactic formalism is its *expressive power*, that is, what events are describable by formulas in the language.

Using a more expressive formalism has both advantages and disadvantages. The advantages are well illustrated by considering Feinberg’s [2000] work on characterizing the common prior assumption (CPA) syntactically. Feinberg considers a language that has knowledge operators and, in addition, belief operators of the form $p_i^\alpha(f)$, which is interpreted as “according to agent i , the probability of f is at least α ”. Feinberg does not have common knowledge in his language, nor does his syntax allow expressions of the form $1/2p_i(f) + 2/3p_i(f') \geq 1$, which can be interpreted as “according to agent i , 1/2 of the probability of f plus 2/3 of the probability of f' is at least 1” or, more generally, expressions of the form $\alpha_1p_i(f_1) + \dots \alpha_kp_i(f_k) \geq \beta$. Expressions of the latter form are allowed, for example, in [Fagin, Halpern, and Megiddo 1990; Fagin and Halpern 1994]. The fact that he does not allow common knowledge causes some technical problems for Feinberg (which he circumvents in an elegant way). Linear combinations of probability terms allow us to make certain statements about expectations of random variables (at least, in the case of a finite state space). Feinberg has an elegant semantic characterization of CPA in the finite case: Roughly speaking, he shows that CPA holds if and only if it is not the case that there

is a random variable X whose expectation agent 1 judges to be positive and agent 2 judges to be negative. Since Feinberg cannot express expectations in his language, he has to work hard to find a syntactic characterization of CPA. With a richer language, it would be straightforward.

This is not meant to be a criticism of Feinberg’s results. Rather, it points out one of the features of the syntactic approach: the issue of the exact choice of syntax plays an important role. There is nothing intrinsic in the syntactic approach that prevents us from expressing notions like common knowledge and expectation. Rather, just as in the semantic approach, where the modeler must decide exactly what the state space should be, in the syntactic approach, the modeler must decide on the choice of language. By choosing a weaker language, certain notions become inexpressible.

Why would we want to choose a weaker language? There are three obvious reasons. One is aesthetic: Just as researchers judge theories on their elegance, we have a notion of elegance for languages. We expect the syntax to be natural (although admittedly “naturalness”, like “elegance”, is in the eye of the beholder) and to avoid awkward expressions. A second is more practical: to the extent that we are interested in doing inference, simpler languages typically admit simpler inference procedures. This intuition can be made formal. There are numerous results characterizing the difficulty of determining whether a formula in a given language is *valid* (that is, true in every state in every structure) [Fagin, Halpern, Moses, and Vardi 1995, Chapter 3]. These results demonstrate that more complicated languages do in fact often lead to more complex decision procedures. However, there are also results showing that, in this sense, there is no cost to adding certain constructs to the language. For example, results of [Fagin, Halpern, and Megiddo 1990; Fagin and Halpern 1994] show that allowing linear combinations of probability terms does not increase the complexity over just allowing operators such as p_i^α . The third reason is that a weaker language might give us a more appropriate level at which to study what we are most interested in. For example, if we are interested only in qualitative beliefs (A is more likely than B), we may not want to “clutter up” the language with quantitative beliefs; that is, we may not want to be able to talk about the exact probability of A and B . A simpler language lets us focus on the key issues being examined. Indeed, roughly speaking, we can think of different languages as providing us with different notions of “sameness”; in other words, “isomorphism” is language-dependent, so any result that holds “up to isomorphism” is really

a language-dependent result.⁵

One other issue: the focus in the semantic approach has been on finding operators that act on events. Not all syntactic expressions correspond in a natural way to operators. For example, recall that in [Fagin, Halpern, and Megiddo 1990], expressions such as $\alpha_1 p_i(f_1) + \cdots + \alpha_k p_i(f_k) \geq \beta$ are considered. This could be viewed as corresponding to a k -ary operator $B^{\alpha_1, \dots, \alpha_k, \beta}$ such that $\omega \in B^{\alpha_1, \dots, \alpha_k, \beta}(E_1, \dots, E_k)$ if $\alpha_1 p_i(E_1) + \cdots + \alpha_k p_i(E_k) \geq \beta$, however, this is not such a natural operator. It seems more natural to associate an event with this syntactic expression directly in terms of a probability distribution on worlds, without bothering to introduce such operators.

Models One of the advantages of the syntactic approach is that it allows us to talk about properties without specifying a model. This is particularly important both if we do not have a fixed model in mind (for example, when discussing rationality, we may be interested in general properties of rationality, independent of a particular model) or when we do have a model in mind, but we do not wish to (or cannot) specify it completely. In general, a formula may be true in many models, not just one.

To the extent that models for formulas have been considered in the economics literature, the focus has been on one special model, the so-called *canonical model* [Aumann 1989]. This canonical model has the property that if a (possibly infinite) collection of formulas is true at some world in some model, then it is true at some world of the canonical model. Thus, in a sense, the canonical model can be thought of as including all possible models.

While the canonical model is useful for various constructions, it also has certain disadvantages. For one thing, it is uncountable. A related problem is that it is far too complicated to be useful as a tool for specifying simple situations. We can think of a formula φ as specifying a collection of structure-world pairs, namely, all the pairs (M, w) such that $(M, w) \models \varphi$. It is well known that, at least in the case of epistemic logic, if a formula φ is satisfiable at all (that is, if $(M, w) \models \varphi$ for some (M, w)), then there is a finite structure M and a world w in M such that $(M, w) \models \varphi$. Instead of focusing on an uncountable canonical structure, it is often much easier to focus on a finite structure.

Thinking in terms of the set of models that satisfy a formula (rather than just one canonical model) leads us to consider a number of different issues.

⁵Thanks to Adam Bjoernsdahl for stressing this point.

In many cases we do not want to construct a model at all. Instead, we are interested in the logical consequences of some properties of a model. A formula φ is a *logical consequence* of a collection Ψ of formulas if, whenever $(M, w) \models \psi$ for every formula $\psi \in \Psi$, then $(M, w) \models \varphi$. For example, we may want to know if some properties of rationality are logical consequences of other properties of rationality; again, this is something that is best expressed syntactically.

4 Conclusion

The point of this short note should be obvious: there are times when syntax is useful, despite its overhead. Moreover, like Molière’s M. Jourdain, who discovers he has been speaking prose all his life [Molière 2013], game theorists have often used syntax without realizing it (or, at least, without acknowledging its use explicitly). However, they have not always taken full advantage of it.

That said, as I have pointed in [Halpern 1999], there are times when semantics is useful, and the overhead of syntax is not worth it. Semantic proofs of validity are typically far easier to carry out than their syntactic counterparts (as anyone who has tried to prove the validity of even a simple formula like $K(\varphi \wedge \psi) \equiv K\varphi \wedge K\psi$ from the axioms knows). I believe that the finer expressive power of syntax is particularly useful when dealing with notions of awareness, and trying to capture how agents react differently to two different representations of the same event. This is something that arguably cannot be done by a semantic approach without essentially incorporating the syntax in the semantics. But when considering an approach where these issues do not arise, a semantic approach may be the way to go. The bottom line here is that it is useful to have both syntactic and semantic approaches in one’s toolkit!

Acknowledgments: Thanks to Adam Bjorndahl for and the two reviewers of this paper for very useful comments.

References

Aumann, R. J. (1989). Notes on interactive epistemology. Cowles Foundation for Research in Economics working paper.

Aumann, R. J. (1995). Backwards induction and common knowledge of rationality. *Games and Economic Behavior* 8, 6–19.

Aumann, R. J. (1999). Interactive epistemology I: knowledge. *International Journal of Game Theory* 28(3), 263–300.

Bacharach, M. (1985). Some extensions of a claim of Aumann in an axiomatic model of knowledge. *Journal of Economic Theory* 37, 167–190.

Balkenborg, D. and E. Winter (1997). A necessary and sufficient epistemic condition for playing backward induction. *Journal of Mathematical Economics* 27, 325–345.

Binmore, K. (2009). *Rational Decisions*. Princeton, N. J.: Princeton University Press.

Bjorndahl, A., J. Y. Halpern, and R. Pass (2013). Language-based games. In *Theoretical Aspects of Rationality and Knowledge: Proc. Fourteenth Conference (TARK 2013)*, pp. 39–48.

Blume, L. and D. Easley (2008). Rationality. In L. Blume and S. Durlauf (Eds.), *The New Palgrave: A Dictionary of Economics*. New York: Palgrave Macmillan.

Blume, L., D. Easley, and J. Y. Halpern (2006). Redoing the foundations of decision theory. In *Principles of Knowledge Representation and Reasoning: Proc. Tenth International Conference (KR '06)*, pp. 14–24. A longer version, entitled “Constructive decision theory”, can be found at <http://www.cs.cornell.edu/home/halpern/papers/behfinal.pdf>.

Fagin, R., J. Geanakoplos, J. Y. Halpern, and M. Y. Vardi (1992). The expressive power of the hierarchical approach to modeling knowledge and common knowledge. In *Theoretical Aspects of Reasoning about Knowledge: Proc. Fourth Conference*, pp. 229–244.

Fagin, R. and J. Y. Halpern (1988). Belief, awareness, and limited reasoning. *Artificial Intelligence* 34, 39–76.

Fagin, R. and J. Y. Halpern (1994). Reasoning about knowledge and probability. *Journal of the ACM* 41(2), 340–367.

Fagin, R., J. Y. Halpern, and N. Megiddo (1990). A logic for reasoning about probabilities. *Information and Computation* 87(1/2), 78–128.

Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi (1995). *Reasoning About Knowledge*. Cambridge, Mass.: MIT Press. A slightly revised paperback version was published in 2003.

Feinberg, Y. (2000). Characterizing common priors in the form of posteriors. *Journal of Economic Theory* 91, 127–179.

Geanakoplos, J. (1989). Game theory without partitions, and applications to speculation and consensus. Cowles Foundation Discussion Paper #914, Yale University.

Gilboa, I. (2010). Questions in decision theory. *Annual Reviews in Economics* 2, 1–19.

Halpern, J. Y. (1999). Set-theoretic completeness for epistemic and conditional logic. *Annals of Mathematics and Artificial Intelligence* 26, 1–27.

Halpern, J. Y. (2003). *Reasoning About Uncertainty*. Cambridge, Mass.: MIT Press.

Halpern, J. Y. and W. Kets (2012). Ambiguous language and differences in beliefs. In *Principles of Knowledge Representation and Reasoning: Proc. Thirteenth International Conference (KR '12)*, pp. 329–338.

Halpern, J. Y. and L. C. Rêgo (2013). Reasoning about knowledge of unawareness revisited. *Mathematical Social Sciences* 66(2), 73–84.

Heifetz, A. and D. Samet (1998). Typology-free typology of beliefs. *Journal of Economic Theory* 82, 324–341.

Molière (2013). *Le Bourgeois Gentilhomme*. Paris: Gallimard.

Samet, D. (1990). Ignoring ignorance and agreeing to disagree. *International Journal of Game Theory* 52, 190–207.